



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/706,360

11/12/2003

Barbara Liskov

CIS03-17(7429)

6805

58406

7590

06/20/2008

BARRY W. CHAPIN, ESQ.
CHAPIN INTELLECTUAL PROPERTY LAW, LLC
WESTBOROUGH OFFICE PARK
1700 WEST PARK DRIVE, SUITE 280
WESTBOROUGH, MA 01581

EXAMINER

PHAN, TUANKHANH D

ART UNIT

PAPER NUMBER

2163

MAIL DATE

DELIVERY MODE

06/20/2008

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

DETAILED ACTION

Response to Amendment

The amendment, file on 3/13/2008, has been entered and acknowledged by the Examiner. Claims 1-27 are pending in the instant application.

Response to Arguments

Applicant's arguments with respect to claims 1-27 have been considered, but they are not persuasive.

Issue I. Applicant argues that neither the cited text nor any other text of Burbeck et al. actually discloses providing a database of bindings of request identifiers to replicas where each binding is a record having a request identifier, a replica identifier and a binding expiration time, the database associated with a first router of the plurality of routers, as required by Applicants' independent claim 1. Indeed, the phrase "database" does not appear anywhere in Burbeck et al., and nowhere in the cited text does Burbeck et al. even suggest that a database of bindings of request identifiers to replicas is stored. Further, nothing in the cited text teaches or suggests a binding expiration time. Indeed, the Examiner seems to attempt to use a service level agreement of a storage service provider as evidence of such an expiration time, see Office Action page 3, but Applicants respectfully submit that such an agreement has absolutely nothing whatsoever to do with a binding of a request identifier to a replica as disclosed throughout Applicants' claims and specification. Finally, Applicants respectfully submit that the Examiner's other support for the teaching or suggestion of this limitation with Burbeck et al., namely col. 3 line 25, has nothing to do with a binding

expiration time, but rather teaches a configurable time interval for sending a request for data if certain conditions are not.

Response I. The Examiner would like to point out that from the reference, Burbeck discloses that if a node that is binding the persistent identifier with servicing server up fails to indicate a connection/reconnection on the expiration of time, the tracking paths and the connection for the data access is maintained by the changing records. Plus, accessing data by a network participant with a servicing server is inherent, for data packets are accessed from one or more databases. Thus, Applicant's argument is not persuasive.

Issue II. Applicants' independent claim 1 also requires maintaining a change log of records entered into the database, each change log entry having a change event generated by the first router and an event number sequential to an event number of a preceding change event in the change log. The Examiner cited to col. 11 lines 15-26 of Burbeck et al. as teaching or suggesting this limitation. However, neither the cited text nor any other text of Burbeck et al. teaches or suggests this limitation. As discussed above with regards to the providing limitation, Burbeck et al. does not teach or suggest a database of bindings of request identifiers to replicas, and thus Burbeck et al. cannot suggest that there are any records entered into such a database, or that a change log of those records is maintained. Entirely separate from this matter, however, the Gossip Monger as cited by the Examiner deals with reputation information. The limitation requires that a change log be maintained of records entered into the database; in other words, the records are, according to the providing limitation, the request identifier (which

the Examiner previously equated to the persistent node identifier of Burbeck et al.), the replica identifier, and the binding expiration time. However, the Gossip Monger does not deal with any of these things but rather deals with reputation information. Further, it is unclear to Applicants, particularly from the cited text, where the Gossip Monger includes a maintained change log of anything. Indeed, the cited text indicates that the Gossip Monger revises node reputations, see col. 11, line 19, but says nothing about keeping a log of such revisions. Thus, for at least any of the reasons given above, Burbeck et al. does not teach or suggest Applicants' independent claim 1, and therefore Applicants' independent claim 1 is allowable over Burbeck et al., either alone or in combination with Srivastava.

Response II. Gossip Monger of Burbeck not only deal with reputation information where network participant records is tracked with the change events, but also with the persistent identifier that is assigned to each node. Plus, status of any event change is tracked so that persistency of connection among network participants and servicing servers is preserved. Thus, Applicant's argument is moot.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1-27 are rejected under 35 U.S.C. 103(a) as being unpatentable over Burbeck et al. (US Pat. 7,143,139), hereinafter Burbeck, and further in view of Srivastava (US Pat. 7,047,315).

Regarding claims 1, 21 and 24-26, Burbeck discloses a method in a computerized device for maintaining a client session in a network having a plurality of routers (i.e. **having servlets called routers**, col. 9, ll. 1-5), the network having an application executed at a plurality of replicas (i.e. **computer program product running a plurality of servlets or routers to receive inbound message**, col. 9, ll. 1-10), comprising the steps of:

providing a database of bindings of request identifiers (i.e. **binding or selecting a persistent node identifier and assigned to each network participant so that the node identifier can be identified and connected to the associated database after it leaves and re-enters the network**, col. 4, ll. 60-67) to replicas where each binding is a record having a request identifier (**a persistent node identifier**, abstract), a replica identifier and a binding expiration time (i.e. **within a configurable time interval allowed**, col. 3, line 25; **a typical service level agreement that specified in the SSP's service commitments included service cycle and storage provide**, col. 5, ll. 45-55), the database associated with a first router of the plurality of routers (col. 9, ll. 1-5; **data base and web-service associated with the storage**, col. 5, ll. 45-55; col. 4, ll. 1-4);

maintaining a change log of records entered into the database, each change log entry having a change event generated by the first router and an event number

sequential to an event number of a preceding change event in the change log (i.e. **the gossip monger manages and stores reputation information as meta-data to update any change overtime include routing paths and events of each node and associated data**, col. 11, ll. 15-26);

maintaining a current version vector associated with the database and the change log, the current version vector entry for the first router being a most recent event number from the change log (**the gossip monger manages and stores reputation information as meta-data to update any change overtime include routing paths and events of each node and associated data**, col. 11, ll. 15-26), the current version vector entry for each other router being a most recent event number received at the first router from that other router (col. 11, ll. 15-26); receiving an update of change events generated at another router in the plurality (i.e. **routing paths and events/interaction history of nodes having medata-data includes current vector entry and refresh/update information gathered**, col. 11, lines 15-26); reconciling the current version vector according to the received update (col. 3, ll. 40-48); and reconciling the database according to the received update such that the client session is maintained (abstract).

While Burbeck concentrates heavily on peer-to-peer network with a router and server, the lack of plurality of servers could be combined. Srivastava discloses a plurality of routers (abstract) that also allow client identifiers to be binded with server/router identifiers (replicas) (col. 4, ll. 59-63). It would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate the

plurality of routers/replicas mapping taught by Srivastava into the above teachings of Burbeck in order to maintain a consistent flow of data and provide a fast-switched over a path to the intended replica so that network overhead and bandwidth are reduced (col. 6, ll. 50-55).

Regarding claim 2, the method of claim 1, Burbeck further discloses wherein the request identifier is a client identifier and an application identifier (i.e. **identifiers are defined for nodes/routers to be identified across network sessions**, abstract).

Regarding claim 3, Burbeck further discloses the method of claim 2 wherein the client identifier is an Internet Protocol address (col. 1, ll. 38-56).

Regarding claim 4, Burbeck further discloses the method of claim 2 wherein the client identifier is a dproxy Internet Protocol address such that the binding associates a dproxy with a replica (Figure 17A; col. 1, ll. 38-56).

Regarding claim 5, the method of claim 1, Burbeck further discloses wherein the step of reconciling the current version vector comprises the steps of: comparing a least recent event number of the router that generated the update to the event number in the current version vector entry, for that router (i.e. **updating and storing content traversal path definition as well as reputation information as appropriate**, col. 23, ll. 30-55); if the least recent event number is in series with the event numbers in the database as determined by the current version vector entry for that other router, then entering the most recent event number of the received update into the current version vector (**the gossip monger manages and stores reputation information as meta-data to update any change overtime include routing paths and events of each**

node and associated data, col. 11, ll. 15-26) entry for the router that generated the update of change events (col. 23, ll. 30-65); and if the least recent event number in the update is not in succession to the event number in the current version vector entry for the router that generated the update of change events, then discarding the received update (col. 23, ll. 30-65).

Regarding claim 6, the method of claim 5 Burbeck further discloses wherein the step of reconciling the database further comprises: if the update was not discarded in the step of reconciling the current version vector, then for each entry of the received update, a) determining whether the received entry has expired (col. 3, ll. 5-45); b) if the received entry has expired, then discarding the entry (col. 3, ll. 5-32); c) if the received entry has not expired, then comparing the request identifier of the received entry with the request identifier in the entries in the database (i.e. **encompassing the response to node's version of its reputation before the failure to satisfy the content request if the time interval expires**, col. 3, ll. 5-32); d) if a matching entry is not found for the received entry, adding the received entry to the database (col. 3, ll. 5-32); e) if a matching entry is found for the received entry, then comparing the application identifier of the received entry with the application identifier of the matching entry (col. 3, ll. 5-32); f) if the application identifiers match, then retaining the entry having a later expiration time in the database (col. 30); and g) if the application identifiers do not match, then retaining an entry selected based on a deterministic function applied to a portion of each entry (abstract; col. 3, lines 5-52, col. 28).

Regarding claim 7, Burbeck further discloses the method of claim 6 wherein the step of retaining an entry based on a deterministic function comprises the steps of applying a function to the application identifiers (col. 23, ll. 30-65); and selecting an entry based on the outcome of the function (col. 23, ll. 30-65).

Regarding claim 8, the method of claim 6, Burbeck further discloses wherein the step of retaining an entry based on a deterministic function comprises the steps of applying the deterministic function to the request identifier (abstract, mapping applications of client to the server) and selecting an entry based on the outcome of the deterministic function (col. 23, lines 30-62).

Regarding claim 9, the method of claim 1, Burbeck further discloses comprising the step of deleting a binding from the database when the expiration time for the binding has been exceeded (col. 23, ll. 30-62).

Regarding claim 10, the method of claim 1 and Burbeck further discloses comprising the step of sending a request for an update of change events to another router in the plurality (col. 23, ll. 30-65); and the step of receiving the update further comprises receiving the update in response to the request (col. 23, ll. 30-65).

Regarding claim 11, the method of claim 1 and Srivastava further discloses further comprising the steps of: periodically generating a first router update of change events (col. 14, ll. 30-49, col. 31); and, transmitting the first router update of change events to at least one other router in the plurality (col. 14, lines 30-49; col. 31).

Regarding claim 12, the method of claim 1 and Srivastava further discloses comprising the steps of: affirming that an update has been received from each router of

the plurality within a predetermined period for each router (col. 14, ll. 30-49; col. 31); if an update has not been received from a router within the predetermined period for that router, requesting an update of change events from that router (col. 14, lines 30-49; col. 31); and if an update is received in response to the request, reconciling the current version vector according to the received update (col. 14, lines 30-49; col. 31); and reconciling the database according to the received update (col. 14, lines 30-49; col. 31).

Regarding claim 13, the method of claim 5 and Srivastava further discloses wherein the step of reconciling the database further comprises the steps of: determining from the received update whether the database has a complete record of changes based on the current version vector (col. 14, lines 30-49; col. 31); if the database does not have a complete record of changes, requesting a replacement database from a router of the plurality of routers (col. 14, lines 30-49; col. 31).

Regarding claim 14, the method of claim 1, Srivastava further discloses, further comprising the step of transmitting a copy of the database and the current version vector to another router of the plurality of routers in response to a request from the other router (col. 14, lines 30-49; col. 31)

Regarding claim 15, see the discussion of claim 1 and Burbeck further discloses wherein the computerized device fails temporarily and recovers, the method further comprising the steps of: writing the first router change log to a persistent storage device (col. 23, ll. 30-60); sending an update of change events written to the change log in the persistent storage device to other routers in the network (col. 23, ll. 40-60); after recovering from failure, requesting a database and an associated version vector from

one of the routers in the plurality (see discussion of claim 1); col. 31); retaining the received database and associated version vector (abstract); reconciling the received database with the change log from the persistent storage device (col. 13, ll. 25-45; col. ll. 30-49); and updating the received version vector (col. 13, ll. 25-45).

Regarding claim 16, the method of claim 1 and Burbeck discloses wherein the received update includes a version vector and the method of maintaining a current version vector further comprises the step of maintaining the current version vector in a version vector table including past version vectors (col. 13, ll. 25-45).

Regarding claim 17, Burbeck discloses the method of claim 16 further comprising the steps of: determining from the version vector table whether the database is current based on the version vector table (col. 13, lines 25-45; col. 23, ll. 40-60); and if the database is not current, then requesting missed change events from a second router in the network (col. 28, lines 1-10).

Regarding claim 18, Burbeck discloses the method of claim 17 wherein each router caches updates received from other routers in the plurality, the method further comprising the step of (abstract): if the router that generated the received update does not respond to the request for missed change events, requesting the missed change events from a second router of the plurality and reconciling the change events into the database and current version vector. It is inherent that since record is periodically ingress update in the table and storage, any miss change of events will be updated entirely upon next update event (col. 23, ll. 40-60).

Regarding claim 19, the method of claim 1 and Burbeck discloses wherein the computerized device fails temporarily and recovers, wherein the step of maintaining a current version vector further comprises the steps of: creating an epoch timestamp from a clock of the computerized device to mark a recovery period (col. 12, ll. 5-20); entering a value pair to the current version vector for the first router, the value pair being an event number and the epoch timestamp (col. 12, ll. 5-35); and the method further comprising the step of after recovery, requesting a database copy and associated version vector from one of the other routers in the plurality (see discussion of claim 1).

Regarding claim 20, Burbeck further discloses the method of claim 19 further comprising the steps of: determining whether a pre-selected time period has passed (col. 3, 25-35); and deleting value pairs before a most recent value pair from the current version vector having timestamps created before the pre-selected time period (col. 11, ll. 15-40).

Regarding claim 22, Burbeck discloses the system of claim 21 wherein the network interface is configured to receive an update of change events for the database from another router in the plurality of routers in the network; and the controller is further configured to reconcile the database according to the received update and to update the current version vector in response to reconciling the database (col. 11, ll. 15-40).

Regarding claim 23, Burbeck discloses the system of claim 21 wherein the controller is configured to transmit periodically, to at least one of the other routers in the plurality, an update of change events and the current version vector (col. 11, ll. 15-26).

Regarding claim 27, Burbeck discloses the method of claim 26 wherein the router is a DNS server and wherein the application identifier in the request is a domain name and wherein the step of routing comprises mapping the request to an Internet Protocol address of the one replica (abstract).

Conclusion

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to TUAN-KHANH PHAN whose telephone number is (571)270-3047. The examiner can normally be reached on 4/5/9.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Don Wong can be reached on 571-272-1834. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2163

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

TKP

/Hung T Vy/
Primary Examiner, Art Unit 2163